

Remarks

Reconsideration of the present application, as amended, is respectfully requested.

Claim 1 has been amended to more clearly and distinctly set forth the present invention. Claims 1-26 remain pending in the application.

Claims 1-6, 8-13 and 15-22 stand rejected under 35 U.S.C. §103(a) as being unpatentable over U.S. Patent Application No. 2003/0074449 to Smith *et al.* (Smith) in view of U.S. Patent Application No. 2002/0004842 to Ghose *et al.* (Ghose) and U.S. Patent Application No. 2002/0138611 A1 to Roe *et al.* (Roe). Claims 7, 14, 23, 24, and 25 stand rejected under 35 U.S.C. §103(a) as being unpatentable over the combination of Smith in view of Ghose, Roe and U.S. Patent No. 5,745,685 to Kirchner *et al.* (Kirchner).

The present invention provides an error recovery mechanism from temporary glitches between two SONET Network Elements (NE) that transport Fibre Channel (FC) traffic over SONET and that use buffer-to-buffer credits. With the present invention, a mechanism keeps track of total number of bytes in transit between the transmitting port and the receive port and ascertains the total available buffer capacity. This mechanism allows the system to recover from temporary glitches and to use the available link bandwidth more efficiently. More specifically, the present invention provides a mechanism that keeps track of the total number of bytes in transit between two SONET ports as well as a robust buffer capacity reporting mechanism which helps the near side self correct its view of the far side buffer capacity during a glitch.

With the present invention, the frame based credit system in FC is replaced with a byte based credit system. This replacement allows the system to keep track of the total buffer available buffer capacity at a much finer granularity to maximize the utilization of the available bandwidth. In each TGFP frame, the receive port, or far side, reports the total number of bytes drained from the FIFO as well as the total available capacity of the FIFO in bytes. The far side keeps track of FIFO capacity by keeping track of the number of bytes written into the far side FIFO. This information is relayed from the far side to the transmitting port in an extension header of the TGFP frame.

In addition to managing the buffer to achieve greater efficiency, the present invention also tracks the total number of bytes in transit between the transmitting port and the receive port.

In order to ascertain the total number of bytes in transit between the transmitting port and receive port, the transmitting port sends a tagged Sequence ID in the linear extension header of at least one TGFP frame. When the Sequence ID is sent to the far side, a counter is started to count the total number of bytes being transmitted. When the tagged Sequence ID is received on the far side, it returns a message back along with the tagged Sequence ID. When the tagged Sequence ID is received, the transmitting port calculates how many bytes are in transit between transmitting port and receive port. If the tagged Sequence ID is dropped in either the transmit or return path, a watchdog timer gets triggered and a new tagged Sequence ID is sent out.

To illustrate, assume that the return acknowledgement frame containing the Sequence ID and buffer capacity and bytes read from the far side buffer gets dropped. If the transmitting port only made its decision to send more bytes on the basis of received Sequence ID 's, it might falsely assume that the frames corresponding to certain Sequence ID 's are still in the pipeline and it would not burst more data out to keep the pipeline full. This would reduce the effective bandwidth utilization.

However, in accordance with the present invention, the transmitting port does not rely on Sequence ID alone. Rather, it also keeps track of how many bytes have been written into the buffer and how many bytes have been read out from the buffer on the far side. Therefore, when the glitch is removed, the far side reports its available capacity to transmitting port which adjusts its view of the total available buffer capacity in response thereto. Since transmitting port is aware of the number of bytes in transit and also knows the total available capacity of the far side buffers, it can also calculate how many more bytes it can burst into the SONET pipeline without overflowing the far side's buffers. Thus, with the present invention, the system can recover from glitches in the link very quickly and provide maximum link bandwidth utilization.

DISCUSSION OF THE RELEVANCY OF THE PRIOR ART

In contrast to the present invention, the systems and methods disclosed in the cited references, either alone or in combination, cannot quickly recover from glitches in the link and provide maximum link bandwidth utilization.

Smith teaches a rudimentary form of buffer to buffer flow control in an attempt to improve utilization of bandwidth. As noted on page 2 of the April 19, 2006 office action:

Smith discloses a buffer to buffer flow control that regulates traffic along a link between a transmitter port and a receiver port by controlling the rate at which the transmitter can send data to the receiver (text omitted). The transmitter is able to transmit a frame along a link only if the receiver has indicated it can accept the frame. (Emphasis added)

As noted by the Examiner, Smith operates in a different manner from the present invention because Smith does not teach or suggest a mechanism that can transmit one or more frames when indications are not received. Rather, Smith specifically teaches that when errors are encountered during the transmission process, the sending node must pause until it receives the indications from the receiver that it is ready to accept additional frames. See for example Smith's discussion at paragraphs 155 and 156. Smith does not describe any mechanism for recovering from a lost receiver ready indicator. Nor does Smith suggest that such a mechanism is necessary or desirable.

However, Smith does suggest one solution for dealing with buffer flow control problems that would minimize the need to pause during transmission is described at paragraphs 189-195. The solution comprises two parts.

The first part of Smith's solution is to remove redundant information, such as idle or other primitive sequences, to minimize the amount of client data that must be stored in the buffer. This teaches away from the present invention, which adds information in the form of the Sequence ID to the frames as well as additional information regarding how many bytes have been read out from the buffer on the far side in the reply.

The second part of Smith's solution is to select the size of the buffers so that it is equal to or larger than the number of buffer credits on the link. Smith's contribution to the art is to carefully select the buffer size so that it is possible to guarantee that the buffer will not overflow due to the flow control mechanism. Relying on having enough memory for the buffer is a poor choice in a world where technology and capacity demands the ability to rapidly adapt to changing conditions. Clearly, the Smith solution is not a flexible solution for improving bandwidth utilization.

Further, Smith fails to recognize that a network glitch can cause the loss of one or more of the buffer ready signals from the receiver. Smith does not recognize, nor does he suggest that a robust buffer capacity reporting mechanism would enable the transmitter to self correct its view

of the far side buffer capacity during a glitch. If, and when, a glitch occurs, the buffer to buffer flow control mechanism taught by Smith breaks down and bandwidth utilization is negatively impacted. The mere fact that the receiving node has available buffer capacity does nothing to improve bandwidth utilization so long as the sending node does not know that the buffer is empty. Thus, the only reasonable solution suggested by the teachings of Smith to overcome a network glitch is to periodically re-initialize the credit count at both the sending node and the receiving node. Periodic re-initialization is clearly undesirable because network utilization is negatively impacted while the nodes re-initialize. Not only did Smith not recognize the very problem solved by the present invention, but he clearly teaches away from the present invention by teaching that the buffer should match or exceed the size necessary to hold the maximum number of frames corresponding to the issued credit. Smith's solution is to use buffer flow control for normal operations and to sit and wait when a glitch occurs.

In contrast to the present invention, Smith does not suggest applying an identification tag to at least an initial one of said GFP-encapsulated client data frames. Rather, Smith is only concerned with maintaining the buffer credit count at a value that is below the maximum buffer credit count limit. As noted above, Smith achieves this goal by teaching the use of an appropriately sized buffer. A prior art reference must be considered in its entirety, i.e., as a whole, including portions that would lead away from the claimed invention. *W.L. Gore & Associates, Inc. v. Garlock, Inc.*, 721 F.2d 1540, 220 USPQ 303 (Fed. Cir. 1983), cert. denied, 469 U.S. 851 (1984). There is no suggestion, either explicit or implicit, that providing additional granularity provided by Ghose would solve the problem solved by the present invention as claimed.

When the present invention is considered as a whole, it is clear that Smith, either alone or in combination with Ghose, Roe or any other reference of record, would not render the claimed invention obvious. In determining the differences between the prior art and the claims, the question under 35 U.S.C. 103 is not whether the differences themselves would have been obvious, but whether the claimed invention as a whole would have been obvious. *Stratoflex, Inc. v. Aeroquip Corp.*, 713 F.2d 1530, 218 USPQ 871 (Fed. Cir. 1983); *Schenck v. Nortron Corp.*, 713 F.2d 782, 218 USPQ 698 (Fed. Cir. 1983). Clearly, the cited references do not render the claimed invention, as a whole, obvious.

The present invention further differs from Smith, either alone or in combination with Ghose, Roe or any other reference of record, in that the present invention receives information that indicates the available memory at the receiver – not just a buffer credit value. The information received by the sending node provides available memory in the buffer in terms of bytes; the round trip transit time and an indication that the remote transport interface has received the at least one of said GFP-encapsulated client data frames having the identification tag. With this information, the sending node can determine available buffer space, the number of frames in transit and, importantly, the number of frames that can be sent without overflowing the buffer at the receiving node.

Combining Smith with any of the other reference of record in this application does not and cannot function to recover from a network glitch. None of the cited references, either alone or in combination, suggest tracking individual tagged frames, determining the transit time or the number of frames in transit following the receipt of the acknowledgment that the tagged frame was received as well as determining the number of bytes that have been read out of the buffer on the far side.

It is well settled that the prior art references themselves must suggest combination. To establish a prima facie case of obviousness, three basic criteria must be met. First, there must be some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings. The Smith reference is devoid of any suggestion to solve the problem solved by the present invention. For example, it does not appear that Smith cares about buffer granularity. Rather, he teaches that knowledge of the availability of unused credits is of paramount importance. As such, there is no teaching, suggestion, or motivation that it would be desirable to provide further granularity of the available space in the receiving node's buffer and it is improper to assert that it would have been obvious for one of ordinary skill in the art to combine the Smith reference with Ghose to achieve the remarkable results of the present invention.

To fulfill the second criteria there must be a reasonable expectation of success in combining the references cited by the Examiner. However, even if Smith and Ghose were combined in the manner suggested by the Examiner, there is no mechanism that applies identification tags to at least one frame and then monitors round trip transit time and in-transit

frames. Importantly, even if combined, the sending node must still pause until it receives a receiver ready signal and, if that signal is lost, bandwidth utilization will be significantly reduced. Thus, even if combined, there is no reasonable expectation of success with the combined references.

To fulfill the third and final criteria the prior art reference (or references when combined) must teach or suggest all the claim limitations. The teaching or suggestion to make the claimed combination and the reasonable expectation of success must both be found in the prior art, not in applicant's disclosure. In re Vaeck, 947 F.2d 488, 20 USPQ2d 1438 (Fed. Cir. 1991). Neither Smith nor Ghose, alone or in combination, suggest how to re-synchronize buffer availability with an erroneous buffer credit value maintained by the sending node. Further, neither Smith nor Ghose describe any mechanism for determining the number of frames in transit between the nodes.

The failure of the Smith and Ghose references to suggest the present invention is not solved by adding the teachings of yet another reference to the hotch pot. Specifically, Roe does not teach or suggest determining the number of frames in transit after acknowledgment of receipt of the tagged frame. Further, combining the teachings of Roe with Smith and Ghose does not suggest the present invention because Smith and Ghose are both concerned with managing buffer credit and not transit time. Thus, there is also no teaching, suggestion or motivation to combine Smith and Ghose with Roe. The mere fact that references can be combined or modified does not render the resultant combination obvious unless the prior art also suggests the desirability of the combination. In re Mills, 916 F.2d 680, 16 USPQ2d 1430 (Fed. Cir. 1990)

With respect to Claims 7, 14, 23, 24, and 25 stand rejected under 35 U.S.C. §103(a) as being unpatentable over the combination of Smith in view of Ghose, Roe and U.S. Patent No. 5,745,685 to Kirchner *et al.* (Kirchner), it is acknowledged that timers, and in particular watchdog timers, are notoriously old in the art. A quick review of the United States Patent Office patent database indicates that over 2,900 issued patents refer to the use of a watchdog timer. The mere fact that timers are known in the art does not necessarily mean that it would have been obvious to incorporate the use of a timer in the manner suggested by the present applicants to solve the problem solved by the present invention.

In Kirchner, multiple timers are used to improve reliability in a client server network. When a packet is sent to the client, the timers represent a means for the NSPP server services to ensure delivery of packets when an Ack (acknowledge) Bit is set. To illustrate, after an Ack Bit is requested, a per-service-defined timer is set to wait for the response to the Ack Bit. If one is not received, then the packet can be re-transmitted, up to the number of per-service-defined retries. When the number of retries is exhausted then the packet is discarded. If the service does not indicate a unique service timer or number of retries, then default values are used.

It is well settled, that a statement that modifications of the prior art to meet the claimed invention would have been 'well within the ordinary skill of the art at the time the claimed invention was made' because the references relied upon teach that all aspects of the claimed invention were individually known in the art is not sufficient to establish a prima facie case of obviousness without some objective reason to combine the teachings of the references. *Ex parte Levengood*, 28 USPQ2d 1300 (Bd. Pat. App. & Inter. 1993). See also *In re Kotzab*, 217 F.3d 1365, 1371, 55 USPQ2d 1313, 1318 (Fed. Cir. 2000)

Indeed, even if one of the timers, as taught by Kirchner, were used to determine the length of time to wait for acknowledgment, the combined prior art of Smith and Ghose and Roe together with Kirchner would not achieve the remarkable results of the present invention. Specifically, it is noted that none of the references teach how to use the timer information to determine the far side buffer capacity based on the number of bytes read out of the buffer and how many frames are in transit. This missing piece of the present invention is neither taught nor suggested by the cited references. Further, there is no motivation for one of ordinary skill in the art at the time the present invention was made to add in this missing piece and use the timer information to determine the far side buffer capacity based on the number of bytes read out of the buffer and how many frames are in transit. When the teachings of Kirchner are combined with Smith and Ghose, the system would merely resend frames once the timer expired. Thus, if the transmitting port had overloaded the receive buffer, Kirchner would teach sending yet more frames in response to dropped frames thereby keeping the receive buffer in an overflow condition and delaying the ability of the system to resolve a buffer overload condition.

To establish prima facie obviousness of a claimed invention, all the claim limitations must be taught or suggested by the prior art. *In re Royka*, 490 F.2d 981, 180 USPQ 580 (CCPA

1974). "All words in a claim must be considered in judging the patentability of that claim against the prior art." In re Wilson, 424 F.2d 1382, 1385, 165 USPQ 494, 496 (CCPA 1970). Simply stated, the cited references do not teach, nor do they suggest, the claim limitations of the present invention.

Applicants submit that since each of the independent claims are nonobvious under 35 U.S.C. 103, then any claim depending therefrom is nonobvious. In re Fine, 837 F.2d 1071, 5 USPQ2d 1596 (Fed. Cir. 1988). Accordingly, applicants submit that the claims, as amended, are in condition for allowance.

Conclusion

To reach a proper determination under 35 U.S.C. 103, the examiner must step backward in time and into the shoes worn by the hypothetical "person of ordinary skill in the art" when the invention was unknown and just before it was made. In view of all the factual information, the examiner must then make a determination whether the claimed invention "as a whole" would have been obvious at that time to that person. Knowledge of applicant's disclosure must be put aside in reaching this determination, yet kept in mind in order to determine the "differences," conduct the search and evaluate the "subject matter as a whole" of the invention. The tendency to resort to "hindsight" based upon applicant's disclosure is often difficult to avoid due to the very nature of the examination process. However, impermissible hindsight and application of *per se* rules must be avoided and the legal conclusion must be reached on the basis of the facts gleaned only from the prior art.

Therefore, the applicants respectfully submit that the claims as now amended are patentable and that the case is in condition to be passed to issue. If a telephone conference would in any way expedite the prosecution of the application, the Examiner is asked to call the undersigned at (408) 377-9195.

Respectfully submitted,
Aka Chan LLP

/Craig E. Shinnors/

Craig E. Shinnors
Reg. No. 33,281

Aka Chan LLP
900 Lafayette Street, Suite 710
Santa Clara, CA 95050
Tel: (408) 701-0035
Fax: (408) 608-1599
E-mail: craig@akachanlaw.com
With a copy to gary@akachanlaw.com